

IN THE CLAIMS

Please amend the claims as follows:

1. (Canceled) A method for reliably storing data in a computer system, comprising:
receiving a piece of data to be stored at a storage system:
writing a first copy of the data to the storage system according to a first data redundancy scheme;
writing a second copy of the data to the storage system according to a second data redundancy scheme;
maintaining metadata of the data written to the storage system according to the second data redundancy scheme;
copying the data written to the storage system according to the second data redundancy scheme to a backup storage system, wherein the copying is performed in response to a defined condition being met; and
removing the data written to the storage system according to the second data redundancy scheme after it has been copied to the backup storage system.
2. (Canceled) The method of claim 1 wherein the stored data includes both blocks of user data and attributes such as file names that are associated with collections of user data.

3. (Currently Amended) A method for reliably storing data in a computer system, comprising:
receiving a piece of data to be stored at a storage system;
writing a first copy of the data to the storage system according to a first data redundancy scheme;
writing a second copy of the data to the storage system according to a second data redundancy scheme;
maintaining metadata of the data written to the storage system according to the second data redundancy scheme;
copying the data written to the storage system according to the second data redundancy scheme to a backup storage system, wherein the copying is performed in response to a defined condition being met; and
removing the data written to the storage system according to the second data redundancy scheme after it has been copied to the backup storage system,
wherein the stored data includes both blocks of user data and attributes such as file names that are associated with collections of user data, and

~~The method of claim 2~~ wherein the metadata is maintained by using a skeleton file system that replicates all the namespace and attribute information of the stored data and a set of delta files each of which records all the data updates that have been made to a corresponding file and that have not been copied to the backup storage system.

4. (Canceled) The method of claim 1 wherein the maintained metadata comprises address information associated with where the data is stored in the storage system according to the first data redundancy scheme.
5. (Canceled) The method of claim 1 further comprises recognizing a failure in the storage system which compromises the reliability of the data stored according to the first data redundancy scheme.
6. (Canceled) The method of claim 5 further comprises responding to the failure by rebuilding an accurate copy of the data stored according to the first data redundancy scheme.
7. (Currently Amended) A method for reliably storing data in a computer system, comprising:
receiving a piece of data to be stored at a storage system;
writing a first copy of the data to the storage system according to a first data redundancy scheme;
writing a second copy of the data to the storage system according to a second data redundancy scheme;
maintaining metadata of the data written to the storage system according to the second data redundancy scheme;
copying the data written to the storage system according to the second data redundancy scheme to a backup storage system, wherein the copying is performed in response

to a defined condition being met; and
removing the data written to the storage system according to the second data redundancy
scheme after it has been copied to the backup storage system,
wherein the method further comprises recognizing a failure in the storage system which
compromises the reliability of the data stored according to the first data redundancy
scheme, and

The method of claim 6 wherein rebuilding an accurate copy of the data stored according to the first data redundancy scheme comprises; utilizing the maintained metadata as a blueprint for combining data retrieved from the backup storage system and data remaining in the store system.

8. (Original) The method of claim 7 wherein data remaining in the storage system comprises data stored according to the second data redundancy scheme.

9. (Canceled) The method of claim 6 wherein rebuilding an accurate copy of the data comprises:

retrieving the maintained metadata of the data written to the storage system according to the second data redundancy scheme;
retrieving the data written to the storage system according to the second data redundancy scheme, where the data has not yet been copied to the backup storage system;
retrieving the data written to the backup storage system; and
rebuilding an accurate copy of the data stored according to the first data redundancy scheme, by utilizing the maintained metadata in combination with the retrieved

data written to the storage system according to the second data redundancy scheme and the retrieved data written to the backup storage system.

10. (Canceled) The method of claim 5 wherein failure comprises a failure of a plurality of physical disks in the storage system.

11. (Canceled) A disaster recovery system, comprising:

a plurality of disks;

a storage controller attached to the disks, wherein the storage controller receives data to be stored on the plurality of physical disks;

first data redundancy scheme logic included within the storage controller for writing a copy of the data received by the storage controller to the physical disks according to a first data redundancy scheme;

second data redundancy scheme logic included within the storage controller for writing a second copy of the data received by the storage controller to the physical disks according to a second data redundancy scheme;

metadata of the data written to the storage system according to the second data redundancy scheme;

data backup logic to copy the data written to the storage system according to the second data redundancy scheme to a backup storage system, wherein the copying is performed in response to a defined condition being met; and

maintenance logic to remove the data written to the storage system according to the second data redundancy scheme after it has been copied to the backup storage

system.

12. (Canceled) The system of claim 11 wherein the stored data includes both blocks of user data and attributes such as file names that are associated with collections of blocks of user data.

13. (Currently Amended) A disaster recovery system, comprising:

a plurality of disks;

a storage controller attached to the disks, wherein the storage controller receives data to be stored on the plurality of physical disks;

first data redundancy scheme logic included within the storage controller for writing a copy of the data received by the storage controller to the physical disks according to a first data redundancy scheme;

second data redundancy scheme logic included within the storage controller for writing a second copy of the data received by the storage controller to the physical disks according to a second data redundancy scheme;

metadata of the data written to the storage system according to the second data redundancy scheme;

data backup logic to copy the data written to the storage system according to the second data redundancy scheme to a backup storage system, wherein the copying is performed in response to a defined condition being met; and

maintenance logic to remove the data written to the storage system according to the second data redundancy scheme after it has been copied to the backup storage

system,
wherein the stored data includes both blocks of user data and attributes such as file names
that are associated with collections of blocks of user data,

~~The method of claim 12~~ wherein the metadata is maintained by using a skeleton file system that replicates all the namespace and attribute information of the stored data and a set of delta files each of which records all the data updates that have been made to a corresponding file and that have not been copied to the backup storage system.

14. (Canceled) The system of claim 11 wherein the maintained metadata comprises address information associated with where the data is stored in the storage system according to the first data redundancy scheme.

15. (Canceled) The system of claim 11 further comprises monitoring logic to recognize a failure in the storage system which compromises the reliability of the data stored according to the first data redundancy scheme.

16. (Canceled) The system of claim 15 further comprises responding to the failure by rebuilding an accurate copy of the data stored according to the first data redundancy scheme.

17. (Currently Amended) A disaster recovery system, comprising:
a plurality of disks;
a storage controller attached to the disks, wherein the storage controller receives data to
be stored on the plurality of physical disks;
first data redundancy scheme logic included within the storage controller for writing a
copy of the data received by the storage controller to the physical disks according
to a first data redundancy scheme;
second data redundancy scheme logic included within the storage controller for writing a
second copy of the data received by the storage controller to the physical disks
according to a second data redundancy scheme;
metadata of the data written to the storage system according to the second data
redundancy scheme;
data backup logic to copy the data written to the storage system according to the second
data redundancy scheme to a backup storage system, wherein the copying is
performed in response to a defined condition being met; and
maintenance logic to remove the data written to the storage system according to the
second data redundancy scheme after it has been copied to the backup storage
system,
wherein the system further comprises monitoring logic to recognize a failure in the
storage system which compromises the reliability of the data stored according to the first
data redundancy scheme,
wherein the system further comprises responding to the failure by rebuilding an accurate
copy of the data stored according to the first data redundancy scheme, and

~~The system of claim 16~~ wherein rebuilding an accurate copy of the data stored according to the first data redundancy scheme comprises, utilizing the maintained metadata as a blueprint for combining data retrieved from the backup storage system and data remaining in the storage system.

18. (Original) The system of claim 17 wherein data remaining in the storage system comprises data stored according to the second data redundancy scheme.

19. (Canceled) The system of claim 16 wherein rebuilding an accurate copy of the data comprises:

retrieving the maintained metadata of the data written to the storage system according to

the second data redundancy scheme;

retrieving the data written to the storage system according to the second data redundancy

scheme, where the data has not yet been copied to the backup storage system;

retrieving the data written to the backup storage system; and

rebuilding an accuracy copy of the data stored according to the first data redundancy

scheme, by utilizing the maintained metadata in combination with the retrieved

data written to the storage system according to the second data redundancy

scheme and the retrieved data written to the backup storage system.

20. (Canceled) The system of claim 15 wherein failure comprises a failure of a plurality of physical disks in the storage system.

21. (Canceled) A computer program product having instruction codes for reliably storing data in a computer system, comprising:

- a set of instruction codes receiving a piece of data to be stored at a storage system;
- a set of instruction codes for writing a first copy of the data to the storage system according to a first data redundancy scheme;
- a set of instruction codes for writing a second copy of the data to the storage system according to a second data redundancy scheme;
- a set of instruction codes for maintaining metadata of the data written to the storage system according to the second data redundancy scheme;
- a set of instruction codes for copying the data written to the storage system according to the second data redundancy scheme to a backup storage system, wherein the copying is performed in response to a defined condition being met; and
- a set of instruction codes for removing the data written to the storage system according to the second data redundancy scheme after it has been copied to the backup storage system.

22. (Canceled) A method for deploying a disaster recovery service in a computer system, comprising:

- integrating computer readable code into a system for receiving a piece of data to be stored at a storage system;
- integrating computer readable code into a system for writing a first copy of the data to the storage system according to a first data redundancy scheme;
- integrating computer readable code into a system for writing a second copy of the data to

the storage system according to a second data redundancy scheme;
integrating computer readable code into a system for maintaining metadata of the data
written to the storage system according to the second data redundancy scheme;
integrating computer readable code into a system for copying the data written to the
storage system according to the second data redundancy scheme to a backup
storage system, wherein the copying is performed in response to a defined
condition being met; and
integrating computer readable code into a system for removing the data written to the
storage system according to the second data redundancy scheme after it has been
copied to the backup storage system.

23. (Canceled) The method of claim 22 wherein the stored data includes both blocks of
user data and attributes such as file names that are associated with collections of blocks
of user data.

24. (Currently Amended) A method for deploying a disaster recovery service in a
computer system, comprising:
integrating computer readable code into a system for receiving a piece of data to be
stored at a storage system;
integrating computer readable code into a system for writing a first copy of the data to the
storage system according to a first data redundancy scheme;
integrating computer readable code into a system for writing a second copy of the data to
the storage system according to a second data redundancy scheme;

integrating computer readable code into a system for maintaining metadata of the data
written to the storage system according to the second data redundancy scheme;
integrating computer readable code into a system for copying the data written to the
storage system according to the second data redundancy scheme to a backup
storage system, wherein the copying is performed in response to a defined
condition being met; and
integrating computer readable code into a system for removing the data written to the
storage system according to the second data redundancy scheme after it has been
copied to the backup storage system,
wherein the stored data includes both blocks of user data and attributes such as file names
that are associated with collections of blocks of user data, and
The method of claim 23 wherein the metadata is maintained by using a skeleton file
system that replicates all the namespace and attribute information of the stored data and a
set of delta files each of which records all the data updates that have been made to a
corresponding file and that have not been copied to the backup storage system.

25. (Canceled) The method of claim 22 wherein the maintained metadata comprises
address information associated with where the data is stored in the storage system
according to the first data redundancy scheme.

26. (Canceled) The method of claim 22 further comprises integrating computer readable
code for recognizing a failure in the storage system which compromises the reliability of
the data stored according to the first data redundancy scheme.

27. (Canceled) The method of claim 26 further comprises integrating computer readable code for responding to the failure by rebuilding an accurate copy of the data stored according to the first data redundancy scheme.

28. (Currently Amended) A method for deploying a disaster recovery service in a computer system, comprising:
integrating computer readable code into a system for receiving a piece of data to be stored at a storage system;
integrating computer readable code into a system for writing a first copy of the data to the storage system according to a first data redundancy scheme;
integrating computer readable code into a system for writing a second copy of the data to the storage system according to a second data redundancy scheme;
integrating computer readable code into a system for maintaining metadata of the data written to the storage system according to the second data redundancy scheme;
integrating computer readable code into a system for copying the data written to the storage system according to the second data redundancy scheme to a backup storage system, wherein the copying is performed in response to a defined condition being met; and
integrating computer readable code into a system for removing the data written to the storage system according to the second data redundancy scheme after it has been copied to the backup storage system.

method further comprising integrating computer readable code for recognizing a failure in the storage system which compromises the reliability of the data stored according to the first data redundancy scheme, and

~~The method of claim 27~~ wherein rebuilding an accurate copy of the data stored according to the first data redundancy scheme comprises; utilizing the maintained metadata as a blueprint for combining data retrieved from the backup storage system and data remaining in the storage system.

29. (Original) The method of claim 28 wherein data remaining in the storage system comprises data stored according to the second data redundancy scheme.

30. (Canceled) The method of claim 27 wherein rebuilding an accurate copy of the data comprises:

retrieving the maintained metadata of the data written to the storage system according to

the second data redundancy scheme;

retrieving the data written to the storage system according to the second data redundancy

scheme, where the data has not yet been copied to the backup storage system;

retrieving the data written to the backup storage system; and

rebuilding an accurate copy of the data stored according to the first data redundancy

scheme, by utilizing the maintained metadata in combination with the retrieved

data written to the storage system according to the second data redundancy

scheme and the retrieved data written to the backup storage system.

31. (Canceled) The method of claim 26 wherein failure comprises a failure of a plurality of physical disks in the storage system.